

Run IIb Event Builder upgrade



Steve Tether
MIT



Ron Rechenmacher
FNAL/CD



Markus Klute
MIT



Bruce Knuteson
MIT

Motivation

Scope

Progress

End game & Installation

Director's Review
Jan 18 2005

Trigger	Rates at $4 \times 10^{32} \text{ cm}^{-2}\text{s}^{-1}$		
	L1	L2	L3
e/μ	2,327 Hz	250 Hz	22 Hz
ν	4,401 Hz	130 Hz	9 Hz
calibration	2,940 Hz	117 Hz	16 Hz
Total	9,668 Hz	497 Hz	47 Hz

Table 6.1: Summary of trigger rates for Higgs search triggers at $\mathcal{L} = 4 \times 10^{32} \text{ cm}^{-2}\text{s}^{-1}$. These triggers are important for many other high- p_T physics analyses.

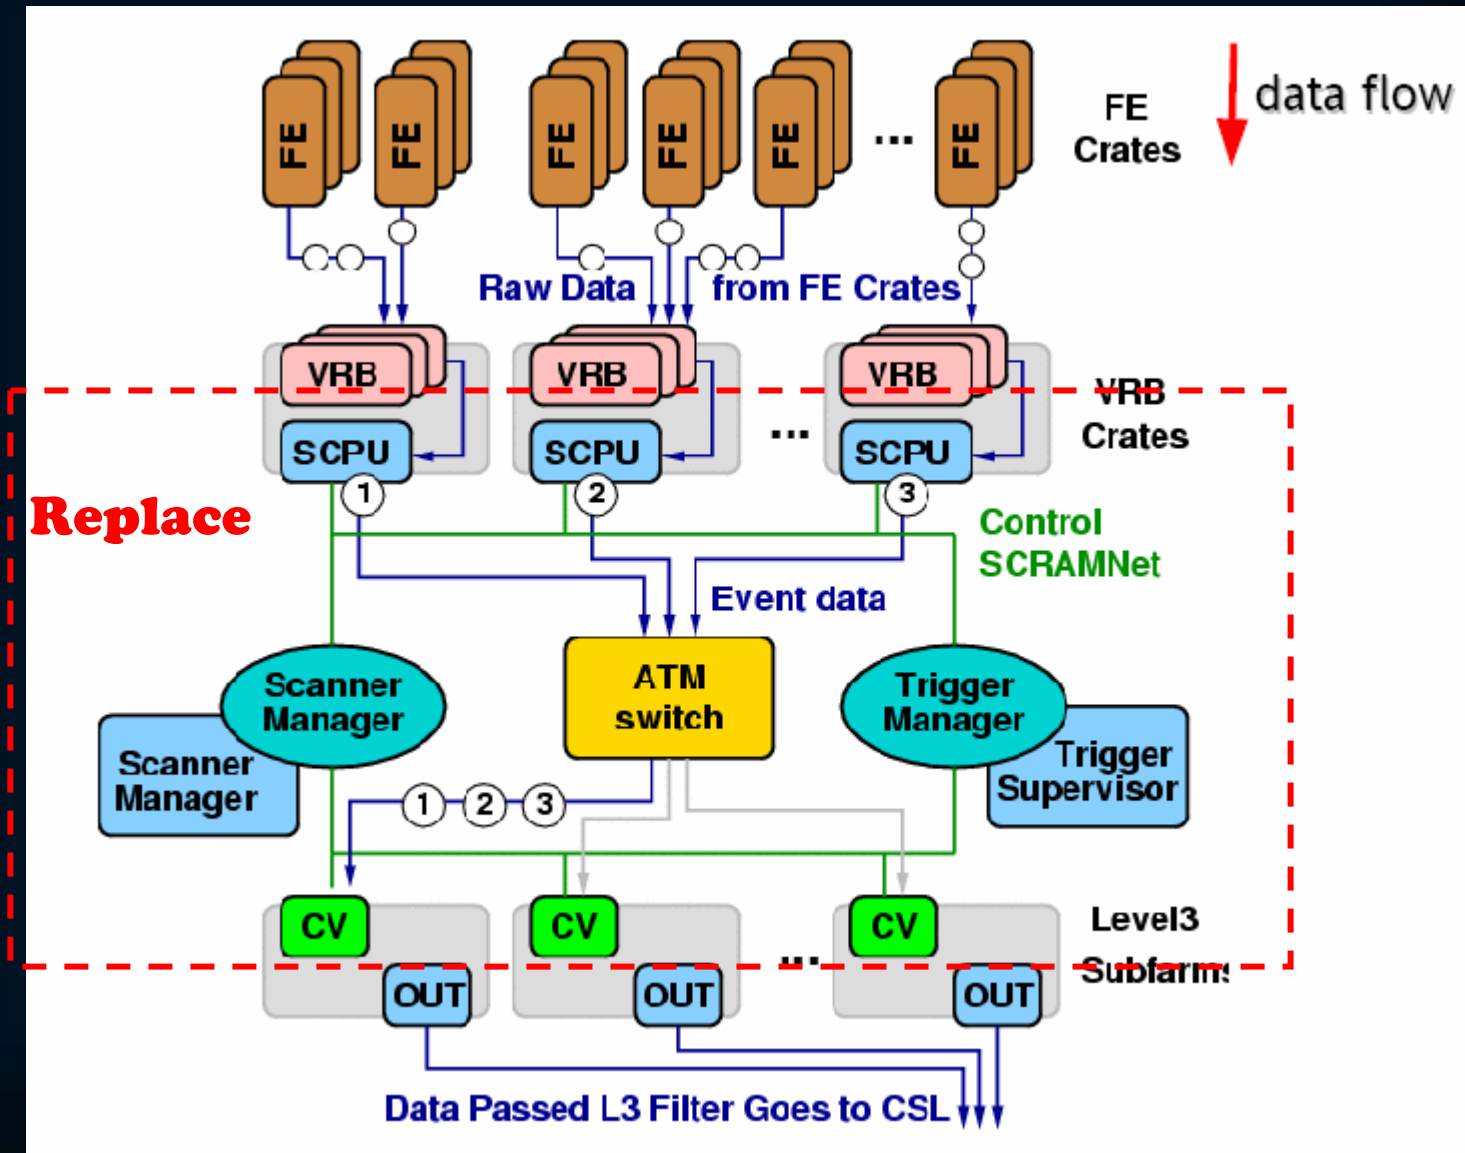
Trigger	$\sigma_{L1}(\text{nb})$	$\sigma_{L2}(\text{nb})$	$\sigma_{L3}(\text{nb})$
High- p_T jets	19,000	60	17
$t\bar{t}$ (all hadronic)	(overlap)	50	5
$\tau\bar{\tau}$	5,000	50	4
$\cancel{E}_T + \tau$	(overlap)	50	4
High- E_T photons	13,500	110	21
di,tri-leptons	1,000	190	45
Total	38,500	660	96
Total rate	15,400 Hz	264 Hz	38 Hz

Table 6.2: Summary of triggers necessary for the CDF Run IIb high- p_T physics program. The estimated rates shown are for an instantaneous luminosity of $\mathcal{L} = 4 \times 10^{32} \text{ cm}^{-2}\text{s}^{-1}$.

Specs

	Runlla	Run llb
Rate:	300Hz	1kHz
Event size:	250kB	500kB
Throughput:	75MB/s	500MB/s

Current Event Builder



Run IIb Event Builder Upgrade

Milestones

- ✓ **Sept 2003** **switch technology choice,
place order for final system switch**
- ✓ **Jan 2004** **arrival of final system switch,
install test stand**

Cisco GbE switch selected and installed

Address http://www.cisco.com/en/US/products/hw/switches/ps708/products_data_sheet09186a008014c49f.html

Google Search Web PageRank 16 blocked AutoFill

CISCO SYSTEMS

Products & Services 60

- PRODUCTS & SERVICES +
- CISCO SWITCHES +
- CISCO CATALYST 6500 SERIES SWITCHES +
- PRODUCT LITERATURE +
- DATA SHEETS +
- Catalyst 6500 Series Gigabit Ethernet Switching Modules

CISCO CATALYST 6500 SERIES SWITCHES

Catalyst 6500 Series Gigabit Ethernet Switching Modules

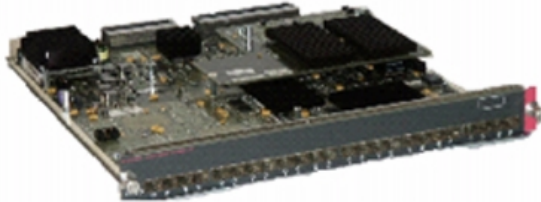
English Select

Data Sheet

Cisco Catalyst 6500 Series Mixed Media Gigabit Ethernet Modules

The Cisco Catalyst® 6500 Series Switch—the premier modular multi-layer switch—provides a wide range of converged services from the wiring closet to the core, to the data center.

Figure 1 Cisco Catalyst 6500 Series Gigabit Ethernet Interface Modules (availability) aCEF daughter card WS-F6700-AFC3



“Complete overkill”

Run IIb Event Builder Upgrade

Milestones

- ✓ **Sept 2003** switch technology choice,
place order for final system switch
- ✓ **Jan 2004** arrival of final system switch,
install test stand
- ✓ **May 2004** SCPU choice,
place order for final system SCPUs,
complete code design
- ✓ **Sept 2004** arrival of final system SCPUs,
code skeleton complete,
start of testing during Summer shutdown

Run IIb Event Builder Upgrade

All hardware now in hand

Gigabit switch (1)
new SCPUs (20)

Additional test stand hardware

16 PCs acting as “Level 3”
Cisco 3750 switch



We are also upgrading Level 3 Converter Nodes (16 nodes)
and Output Nodes (8 nodes)

Purchase order is being placed

Run IIb Event Builder Upgrade

Summer 2004 Shutdown Milestones



Oct 1 2004

Small test

- 1 new SCPU**
- dummy Level 3**
- limited Run Control interface**
- no interface to Trigger Manager**





Nov 1 2004

Medium test

- several new SCPUs**
- dummy Level 3**
- functional Run Control interface**
- functional Trigger Manager interface**

Dec 1 2004

Large test

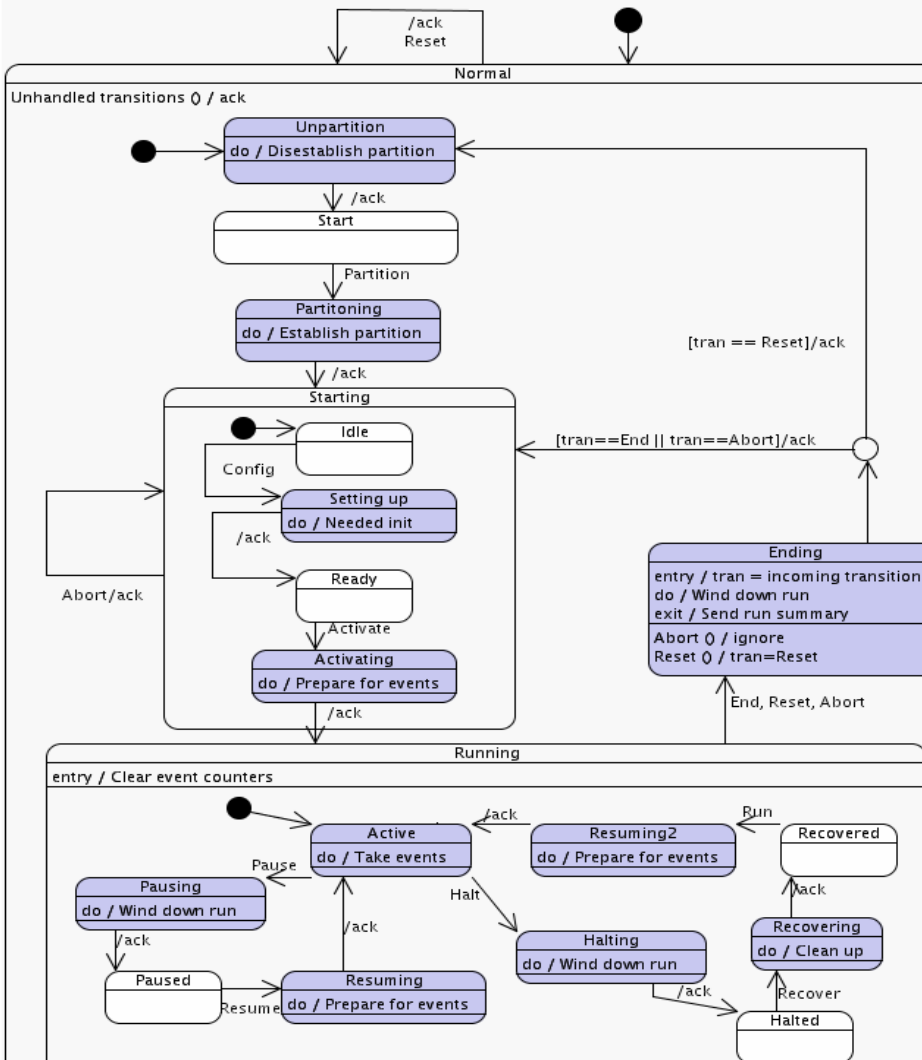
- all SCPUs**
- full Level 3**
-  **full Run Control interface**
-  **full Trigger Manager interface**

Extensive code design

(<http://mit.fnal.gov/~klute/design/>)



Use case: Obey state machine



Real implementation

(<https://www-cdfonline.fnal.gov/cgi-bin/cvsweb.cgi/cdfevb2/>)



emacs@b0evb2gate.fnal.gov

File Edit Options Buffers Tools C++ Help

```
// Markus Klute, klute@fnal.gov
// MIT
// created: Sun May 09, 2004
// $Id: StateMachine.cc,v 1.20 2004/11/17 03:25:28 klute Exp $

#include <scpu/StateMachine.hh>
#include <scpu/StateMachineDriver.hh>
#include <scpu/Messages.hh>
#include <evb/MainTrans.hh>
#include <evb/AceUtil.hh>
#include <evb/Message.hh>
using evb::Message;
#include <boost/shared_ptr.hpp>
using boost::shared_ptr;

namespace scpu {

const QEvent StateMachine::AckSetupVrb_EVENT = { AckSetupVrb_SIG, 0, 0 };
const QEvent StateMachine::AckPrepareL3_EVENT = { AckPrepareL3_SIG, 0, 0 };
const QEvent StateMachine::Configuration_EVENT = { Configuration_SIG, 0, 0 };
const QEvent StateMachine::AckConfigVrb_EVENT = { AckConfigVrb_SIG, 0, 0 };
const QEvent StateMachine::AckConfigRun_EVENT = { AckConfigRun_SIG, 0, 0 };
const QEvent StateMachine::AckVrbReset_EVENT = { AckVrbReset_SIG, 0, 0 };
const QEvent StateMachine::Setup_EVENT = { Setup_SIG, 0, 0 };

// *** Waiting for EvbVrbList.
QSTATE StateMachine::partitioning(QEvent const *e)
{
    switch(e->sig)
    {
        case Q_ENTRY_SIG:
        {
            ACE_DEBUG ((LM_DEBUG, "(%P/%t) StateMachine : partitioning-ENTRY.\n"));
            return 0;
        }
        case Q_EXIT_SIG:
        {
            ACE_DEBUG ((LM_DEBUG, "(%P/%t) StateMachine : partitioning-EXIT.\n"));
            return 0;
        }
        case Done_SIG:
        {
            ACE_DEBUG ((LM_DEBUG, "(%P/%t) StateMachine : partitioning-Done.\n"));
            Q_TRAN(&StateMachine::starting);
            shared_ptr<const Message> msg(new evb::AckPartition());
            driver().send(msg);
        }
    }
}
```

Run IIb Event Builder Upgrade

Milestones

- ✓ **Sept 2003** switch technology choice,
place order for final system switch
- ✓ **Jan 2004** arrival of final system switch,
install test stand
- ✓ **May 2004** SCPU choice,
place order for final system SCPUs,
complete code design
- ✓ **Sept 2004** arrival of final system SCPUs,
code skeleton complete,
start of testing during Summer shutdown
- ✓ **Jan 2005** completion of Shutdown testing
- May 2005** code complete,
system ready for installation
- Sept 2005** Run IIb Event Builder installed

Study of the Event Size and the Performance of the Run IIb Event Builder System.

CDF-7327

Markus Klute


Massachusetts Institute of Technology

December 13, 2004

This note describes an analysis of the event fragment size in Event Builder VME crates and the performance of the Run IIb event builder system. A recommendation is given to increase the number of VME crates in order accommodate the estimated event size at the projected Tevatron peak luminosity.

Our bandwidth limit will be the VME backplane


Achiev[ed/able] rates



$$t = 2 \cdot (15\mu s + (n_{VRB} - 1) * 5\mu s) + d / (50MB/s),$$

(code currently in flux)

Projected Event Size through DAQ (*sans* Silicon) VRB crates



$(300 \text{ kB} / 6 \text{ DAQ crates}) / (30 \text{ MB/sec VME readout}) = 1.7 \text{ ms} \rightarrow 600 \text{ Hz}$

$(300 \text{ kB} / 12 \text{ DAQ crates}) / (30 \text{ MB/sec VME readout}) = 0.8 \text{ ms} \rightarrow 1.2 \text{ kHz}$

We are therefore adding 6 extra DAQ VRB crates

Cost: ≈ 40 k\$

6 VMIC 7805 boards at \$2700 each = 16 k\$

4 Rittal VIPA crates at \$4000 each = 16 k\$

Misc infrastructure (power cables) ~ 10 k\$

We are also rearranging Silicon VRB crates to balance load

Report of the December 17, 2004 Run2b Event Builder Upgrade Review

William Badgett, Guillermo Gomez-Ceballos, Steve Nahn, Jim Patrick, Mel Shochet

Run 2b Event Builder Upgrade Review Committee

Abstract

This is a report on the findings of the Review committee for the Run 2B Event Builder based on the presentations of the [review of December 17, 2004](#).

Contents

1	Introduction	1
2	General	2
3	Comments and Recommendations from the committee	3
3.1	Throughput calculations	3
3.2	Error Handling and Efficiency of Operation	3
3.3	Commissioning and Operations Plan	4
3.4	Additional Crates	5
4	Summary	6

Committee recommendations:

1. Optimize code

We are concentrating on optimization while completing the full functionality of the system. Last week we eliminated two of two time-consuming copies of data in memory, and we are beginning to bundle messages for increased speed.

2. Error handling

We are adding dumps of event information when an error is seen. Our design will handle most errors with a “HRR” (CDF jargon), which costs a few seconds.

Committee recommendations:

3. Commissioning and operations

We have added hardware to the VRB crates so that switching between new and old systems costs ~ 15 min

We will be conducting tests with the CDF DAQ this Spring, using our test stand to ensure good use of this time

New personnel (2 graduate students, 1 postdoc) are being brought on board to learn system, carry pager

4. Additional VRB crates

We are adding 6 additional DAQ VRB crates to the system, and are rearranging Silicon VRBs.

Summary

(1 of 2)

- **All hardware is in hand**
- **Significant progress in code development**
- **Desired rates have been demonstrated**
- **Combination of**
 - **current CDF event builder expertise** (Steve Tether)
 - **existing DO event builder expertise** (Ron Rechenmacher)
 - **young blood** (Markus Klute)**has proven valuable**
- **Thanks to many involved in the CDF DAQ and operations**
Frank, JJ, Bill, Steve, Peter, Pat, Cheng-Ju, Rob, Dervin,
and others

Summary

(2 of 2)

- **We are on target for completion by May-Aug 2005**
- **Possibility exists to install before Summer Shutdown 2005**
- **Quick swap between old and new systems enables commissioning of new system during hours without beam over the next few months**
- **We will use (occasional) access to the CDF DAQ to do this commissioning**

**Looking forward to a tripling of rate
and sextupling of throughput into CDF Level 3**